arXiv:1610.08360v1 [stat.ME] 26 Oct 2016

Efficiently estimating the error distribution in nonparametric regression with responses missing at random

Justin Chown[a*] and Ursula U. Müller[b*]

Abstract. This article considers nonparametric regression models with multivariate covariates and with responses missing at random. We estimate the regression function with a local polynomial smoother. The residual-based empirical distribution function that only uses complete cases, i.e. residuals that can actually be constructed from the data, is shown to be efficient in the sense of Hájek and Le Cam. In the proofs we derive, more generally, the efficient influence function for estimating an arbitrary linear functional of the error distribution; this covers the distribution function as a special case. We also show that the complete case residual-based empirical distribution function admits a functional central limit theorem. The article concludes with a small simulation study investigating the performance of the complete case residual-based empirical distribution function.

keywords: efficient estimator, empirical distribution function, local polynomial smoother, martingale transform, missing at random, nonparametric regression, test for normal errors, transfer principle

2010 AMS Subject Classifications: Primary: 62G05; Secondary: 62G08, 62G20

1. Introduction and main result

An important tool for making decisions about goodness-of-fit and lack-of-fit is the residual-based empirical distribution function. This has been studied in many articles. Stute (1997) and Khmaladze and Koul (2004, 2009), for example, test parametric hypotheses about the regression function in nonparametric models. Neumeyer and Van Keilegom (2010) study additivity tests in heteroskedastic nonparametric regression. Müller, Schick and Wefelmeyer (2012) test for normal errors.

In this article we study the nonparametric regression model

$$Y = r(X) + \varepsilon, \tag{1.1}$$

with the error ε independent of the covariate vector X. Nonparametric models are particularly useful for residual-based inference because residuals constructed from them are usually consistent. We are interested in the case where responses Y are missing, i.e. we observe the sample $(X_1, \delta_1 Y_1, \delta_1), \ldots, (X_n, \delta_n Y_n, \delta_n)$, where δ is an indicator variable which equals one, if Y is observed, and zero, otherwise. In practical applications, most datasets contain missing responses. It is important to choose appropriate statistical methods that ensure conclusions are not biased.

* *Correspondence may be addressed to either author:*
[a] *Fakultät für Mathematik, Lehrstuhl für Stochastik, 44780 Bochum, DE*
Email: justin.chown@ruhr-uni-bochum.de
[b] *Department of Statistics, Texas A&M University, College Station, TX 77843-3143, USA*
Email: uschi@stat.tamu.edu .

We make the assumption that responses are *missing at random* (MAR). This means that the probability that Y is observed depends only on the covariates,

$$P(\delta = 1|X, Y) = P(\delta = 1|X) = \pi(X).$$

We will refer to the model with responses missing at random as the *MAR model*. MAR is a common assumption and is reasonable in many situations (see Little & Rubin, 2002, Chapter 1). As an example, consider missing responses to a survey question about income. If additional data (X) about medical conditions were available, we might see that the response probabilities (π) are smaller for subjects diagnosed with depression. In this case the missing mechanism is ignorable since π depends only on fully observed data X, i.e. it can be estimated from the data. More examples of missing data can be found in Tsiatis (2006), in Liang, Wang and Carroll (2007), in Molenberghs and Kenward (2007), and in Efromovich (2011a, 2011b).

We show in this article that the residual-based empirical distribution function $\hat{\mathbb{F}}_c$ given in equation (1.3) below is an efficient estimator of the unknown error distribution function F. This estimator uses only the complete data pairs (X, Y), i.e. the available residuals $\hat{\varepsilon}_{j,c} = Y_j - \hat{r}_c(X_j)$, where $\hat{r}_c$ is a suitable complete case estimator of the regression function.

Demonstrating this requires two steps. First we show that $\hat{\mathbb{F}}_c$ satisfies the uniform stochastic expansion

$$\sup_{t\in\mathbb{R}} \left| \hat{\mathbb{F}}_c(t) - \frac{1}{N}\sum_{j=1}^{n} \delta_j \mathbf{1}\big[\varepsilon_j \le t\big] - f(t)\frac{1}{N}\sum_{j=1}^{n} \delta_j \varepsilon_j \right| = o_p(n^{-1/2}). \tag{1.2}$$

Here f is the error density and $N = \sum_{j=1}^{n} \delta_j$ is the number of complete cases. Then we show that an estimator of F that admits this expansion is asymptotically efficient in the sense of Hájek and Le Cam. This follows from the arguments in Section 2, where we derive, more generally, the efficient influence function for estimating an arbitrary linear functional $E[h(\varepsilon)]$, which covers $F(t) = E[1(\varepsilon \le t)]$ as a special case. We conclude that an estimator $\hat{\mathbb{F}}_c$ with expansion (1.2) is indeed efficient for F.

We handle part of the proof that (1.2) holds by using the *transfer principle* for complete case statistics in Koul, Schick and Müller (2012). This principle makes it possible to adapt results for the model where all data are fully observed, the *full model*, to missing data models. In particular, we can use the complete case version $\hat{r}_c$ of an estimator $\hat{r}$ in the fully observed data model (i.e. all indicators δ_j are equal to one). Müller, Schick and Wefelmeyer (2009) obtain expansion (1.2) for the full model using a local polynomial smoother to estimate the regression function r, and these authors derive useful results the local polynomial estimator of the regression function that are suitable to the missing data model considered here. See also Neumeyer and Van Keilegom (2010), who consider heteroskedastic nonparametric regression.

In order to summarize the main result by Müller, Schick and Wefelmeyer (2009) (Theorem 1 below), we introduce some notation. Let $i = (i_1, \dots, i_m)$ be a multi-index and write $I(k)$ for the set of multi-indices that satisfy $i_1 + \cdots + i_m \le k$. Müller, Schick and Wefelmeyer (2009) estimate r by a local polynomial smoother $\hat{r}$ of degree d. It is defined as the component $\hat{\beta}_0$ corresponding to the multi-index $0 = (0, \dots, 0)$ of a minimizer

$$\hat{\beta} = \operatorname*{arg\,min}_{\beta=(\beta_i)_{i\in I(d)}} \sum_{j=1}^{n} \left\{ Y_j - \sum_{i\in I(d)} \beta_i \psi_i\left(\frac{X_j - x}{c_n}\right)\right\}^2 w\left(\frac{X_j - x}{c_n}\right),$$

where

$$\psi_i(x) = \frac{x_1^{i_1}}{i_1!} \cdots \frac{x_m^{i_m}}{i_m!}, \qquad x = (x_1, \dots, x_m) \in \mathbb{R}^m,$$

$w(x) = w_1(x_1) \cdots w_m(x_m)$ is a product of densities, and $\{c_n\}_{n\ge1}$ is a bandwidth sequence.

The estimator $\hat{r}$ permits the desired expansion, if the assumptions of Theorem 1 (below) are satisfied. This requires, in particular, the regression function r belongs to the Hölder space $H(d, \gamma)$, i.e. it has continuous partial derivatives of order d (or higher), and that the partial

derivatives of order d are Hölder with exponent γ. The choice of the degree d of the local polynomial smoother will also depend on smoothness and moment conditions on the error density, and on the dimension of the covariate vector. In our simulation study in Section 3, we consider an infinitely differentiable regression function r and a one-dimensional covariate X, which allows us to use a locally linear smoother. Theorem 1 from Müller, Schick and Wefelmeyer (2009) is proved under the following assumption on the covariate distribution:

ASSUMPTION 1. The covariate vector X is quasi-uniform on the cube $[0,\,1]^m$, i.e. X has a density which is bounded and bounded away from zero on $[0,\,1]^m$.

THEOREM 1 (Theorem 1 of Müller, Schick and Wefelmeyer, 2009). *Let Assumption 1 be satisfied. Suppose that the regression function r belongs to $H(d,\gamma)$ with $s = d+\gamma > 3m/2$. Let that the error variable have mean zero, a finite moment of order $\zeta > 4s/(2s-m)$ and a density f that is Hölder with exponent $\xi > m/(2s-m)$. Consider the estimator $\hat{r}$ from above with densities $w_1,\dots,w_m$ that are $(m+2)$-times continuously differentiable with compact support $[-1,\,1]$. Finally, let the bandwidth sequence satisfy $c_n \sim (n\log n)^{-1/(2s)}$. Then, writing $\hat{\varepsilon}_j = Y_j - \hat{r}(X_j)$,*

$$\sup_{t\in\mathbb{R}}\left|\frac{1}{n}\sum_{j=1}^{n}\Big\{\mathbf{1}\big[\hat{\varepsilon}_j\le t\big]-\mathbf{1}\big[\varepsilon_j\le t\big]-\varepsilon_j f(t)\Big\}\right| = o_p(n^{-1/2}).$$

We can apply the transfer principle for asymptotically linear statistics given by Koul, Müller and Schick (2012) to adapt the results from Theorem 1 for the MAR model as follows. The complete case estimator for $F(t)$ is given by

$$\hat{\mathbb{F}}_c(t) = \frac{1}{N}\sum_{j=1}^{n}\delta_j\mathbf{1}\big[\hat{\varepsilon}_{j,c}\le t\big] = \frac{1}{N}\sum_{j=1}^{n}\delta_j\mathbf{1}\big[Y_j-\hat{r}_c(X_j)\le t\big], \tag{1.3}$$

where $\hat{r}_c$ is the complete case version of $\hat{r}$, i.e. $\hat{r}_c$ is given by the component $\hat{\beta}_{c0}$ of a minimizer

$$\hat{\beta}_c = \operatorname*{arg\,min}_{\beta=(\beta_i)_{i\in I(d)}}\sum_{j=1}^{n}\delta_j\bigg\{Y_j-\sum_{i\in I(d)}\beta_i\psi_i\bigg(\frac{X_j-x}{c_n}\bigg)\bigg\}^2 w\bigg(\frac{X_j-x}{c_n}\bigg). \tag{1.4}$$

Using the transfer principle requires the conditional distribution of (X,Y) given $\delta=1$ to meet the assumptions on the (unconditional) joint distribution of (X,Y) from Theorem 1. In our case, it is easy to see how this requirement affects only the covariate distribution G: the MAR assumption combined with the independence of X and ε yield that ε and (X,δ) are independent. This implies the parameters f and r stay the same when switching from the unconditional to the conditional distribution. In particular, the complete case statistic $\hat{\mathbb{F}}_c(t)$ is a consistent estimator for $F(t)$ in the MAR model (since F remains unchanged). Hence, we can keep all but one of our assumptions: only Assumption 1 must be restated.

ASSUMPTION 2. The conditional distribution of the covariate vector X given $\delta=1$ is quasi-uniform on the cube $[0,1]^m$, i.e. it has a density which is bounded and bounded away from zero on $[0,1]^m$.

The transfer principle implies the complete case version of the estimator from Theorem 1 has the corresponding expansion (1.2). This expansion is equivalent to

$$\sup_{t\in\mathbb{R}}\left|\frac{1}{n}\sum_{j=1}^{n}\frac{\delta_j}{E\delta}\Big\{\mathbf{1}\big[\hat{\varepsilon}_{j,c}\le t\big]-\mathbf{1}\big[\varepsilon_j\le t\big]-\varepsilon_j f(t)\Big\}\right| = o_p(n^{-1/2}).$$

Hence, we have, uniformly in $t\in\mathbb{R}$,

$$\hat{\mathbb{F}}_c(t) = \frac{1}{n}\sum_{j=1}^{n}\frac{\delta_j}{E\delta}\mathbf{1}\big[\hat{\varepsilon}_{j,c}\le t\big] + o_p(n^{-1/2}) = F(t) + \frac{1}{n}\sum_{j=1}^{n}b(\delta_j,\varepsilon_j,t) + o_p(n^{-1/2}),$$

where $b(\delta, \varepsilon, t) = (\delta/E\delta)\{\mathbf{1}[\varepsilon \le t] - F(t) + \varepsilon f(t)\}$ is the influence function. This is indeed the *efficient influence function* for estimating $F(t)$: see Corollary 2 in Section 2. This brings us to the main result of this paper.

THEOREM 2. *Consider the nonparametric regression model with responses missing at random. Suppose the assumptions of Theorem 1 are satisfied, with Assumption 2 in place of Assumption 1. Then the* complete case *estimator $\hat{\mathbb{F}}_c$ of the error distribution function F satisfies the stochastic expansion* (1.2)*,*

$$\sup_{t\in\mathbb{R}} \left| \frac{1}{N} \sum_{j=1}^{n} \delta_j \Big\{ \mathbf{1}\big[\hat{\varepsilon}_{j,c} \le t\big] - \mathbf{1}\big[\varepsilon_j \le t\big] - \varepsilon_j f(t) \Big\} \right| = o_p(n^{-1/2}).$$

If the error density f furthermore fulfills Assumption 3, stated in Section 2, then $\hat{\mathbb{F}}_c$ is asymptotically efficient in the sense of Hájek and Le Cam for estimating F, with influence function

$$b(\delta, \varepsilon, t) = \frac{\delta}{E\delta} \Big\{ \mathbf{1}\big[\varepsilon \le t\big] - F(t) + \varepsilon f(t) \Big\}.$$

REMARK 1. If the transfer principle were not available, the expansion in Theorem 2 could be derived by mimicking the (rather elaborate) proofs of Lemma 1 in Müller, Schick and Wefelmeyer (2009) and of Theorem 2.2 in Müller, Schick and Wefelmeyer (2007), who estimate the error distribution in a general semiparametric regression model. The arguments are essentially the same – what is new now is the presence of indicators.

Analogously to Müller, Schick and Wefelmeyer (2009), derive an approximation $\hat{a}_c(x)$ of the difference $\hat{r}_c(x) - r(x)$,

$$\sup_{x\in[0,1]^m} \big| \hat{r}_c(x) - r(x) - \hat{a}_c(x) \big| = o_p(n^{-1/2}); \tag{1.5}$$

see equation (1.4) in that paper. Note, the events $\{\hat{\varepsilon}_{j,c} \le t\}$ and $\{\varepsilon_j \le t + \hat{r}_c(x) - r(x)\}$ are equivalent. Combining this fact and (1.5) with replacing the two empirical distribution functions $\hat{\mathbb{F}}_c$ and $N^{-1}\sum_{j=1}^{n} \delta_j \mathbf{1}[\varepsilon_j \le t]$ in the proof of Theorem 2.2 in Müller, Schick and Wefelmeyer (2007) yields

$$\sup_{t\in\mathbb{R}} \left| \frac{1}{n} \sum_{j=1}^{n} \frac{\delta_j}{E\delta} \Big\{ \mathbf{1}\big[\hat{\varepsilon}_{j,c} \le t\big] - \mathbf{1}\big[\varepsilon_j \le t\big] - F_{\hat{a}_c}(t) - F(t) \Big\} \right| = o_p(n^{-1/2}),$$

writing

$$F_{\hat{a}_c}(t) = E\left[\frac{\delta_j}{E\delta} \mathbf{1}\big[\varepsilon \le t + \hat{a}_c(X)\big] \right] = E\Big[\mathbf{1}\big[\varepsilon \le t + \hat{a}_c(X)\big] \,\Big|\, \delta = 1 \Big] = \int_{[0,1]^m} F\big(t + \hat{a}(x)\big)\, G_1(dx).$$

Here G_1 denotes the conditional distribution of X given $\delta = 1$. A Taylor expansion applied to the difference $F_{\hat{a}_c}(t) - F(t)$ in the above expansion gives

$$\sup_{t\in\mathbb{R}} \left| \frac{1}{n} \sum_{j=1}^{n} \frac{\delta_j}{E\delta} \left\{ \mathbf{1}\big[\hat{\varepsilon}_{j,c} \le t\big] - \mathbf{1}\big[\varepsilon_j \le t\big] - f(t) \int_{[0,1]^m} \hat{a}_c(x)\, G_1(dx) \right\} \right| = o_p(n^{-1/2}).$$

The desired expansion now follows from this combined with

$$\int_{[0,1]^m} \hat{a}_c(x)\, G_1(dx) = \frac{1}{n} \sum_{j=1}^{n} \frac{\delta_j}{E\delta} \varepsilon_j + o_p(n^{-1/2}).$$

The last approximation is the complete case version of equation (1.3) in Müller, Schick and Wefelmeyer (2009). It can be verified by inspecting the proof of Lemma 1 in that paper, where properties of locally polynomial smoothers are derived. Keep in mind that our estimators are constructed from the complete cases (equation (1.4) above), which explains the indicators in the above formula.

Note, the uniform expansion (1.2) implies $\hat{\mathbb{F}}_c$ satisfies a functional central limit theorem, and the efficiency property of the estimator $\hat{\mathbb{F}}_c$ guarantees that competing estimators will not be able to outperform it in large samples. This includes estimators based on imputations that attempt to replace the missing responses. The article is organized as follows. We provide the efficient influence function for estimating linear functionals of the error distribution function F in Section 2, and we specialize these results to estimators of F. In Section 3, we illustrate this result with simulations for two examples. The first example demonstrates the efficiency property of the complete case estimator $\hat{\mathbb{F}}_c$ by comparing it with a 'tuned' estimator using an imputation technique that is in the spirit of González-Manteiga and Pérez-González (2006). For our second example, we perform simulations similar to those in Müller, Schick and Wefelmeyer (2012), who use a martingale transform approach to test for normal errors in the full model. The test statistics involve the estimators from the first example.

2. Efficiency

In this section we provide the *efficient influence function* for estimating the linear functional $E[h(\varepsilon)]$ using observations $(X_i, \delta_i Y_i, \delta_i)$, $i = 1, \ldots, n$. We first follow the arguments of Müller, Schick and Wefelmeyer (2006), who study efficient estimation of general differentiable functionals with data of the above form. We summarize their main arguments and refer to that paper for more details. We then focus on the functional $E[h(\varepsilon)]$, which Müller, Schick and Wefelmeyer (2004) study in the full model. This allows us to adapt parts of their proofs to the MAR model considered here. To begin, we will require the Fisher information for location of the error distribution to be finite:

Assumption 3. The error density f is absolutely continuous with almost everywhere derivative f' satisfying

$$J = \int_{-\infty}^{\infty} \ell^2(z) f(z)\, dz < \infty,$$

where J is the Fisher information for location and $\ell = -f'/f$ is the score function.

We do not assume a parametric model for the regression function or for the distribution of the observations. The parameter set Θ of the statistical model therefore includes a family of covariate distributions $\mathscr{G}$ satisfying Assumption 1, a family of error distributions $\mathscr{F}$ satisfying Assumption 3, a space of regression functions $\mathscr{R}$ that belong to $H(d, \gamma)$, and a family of response probability distributions $\mathscr{B}$ that are characterized by proportion functions mapping $[0,\, 1]^m$ to $(0,\, 1]$. It follows that we can write $\Theta = \mathscr{G} \times \mathscr{F} \times \mathscr{R} \times \mathscr{B}$.

Since the construction of the efficient influence function utilizes the directional information in Θ, we will now identify the set $\dot{\Theta}$ of all perturbations related to the statistical model, which may be thought of as directions. The joint distribution $P(dx, dy, dz)$ depends on the marginal distribution $G(dx)$ of X, the conditional probability $\pi(x)$ that δ equals one given $X = x$, and the conditional distribution $Q(dy \,|\, x)$ of Y given $X = x$:

$$P(dx, dy, dz) = G(dx) B_{\pi(x)}(dz \,|\, x) \Big\{ z Q(dy \,|\, x) + (1 - z)\delta_0(dy) \Big\},$$

where $B_p = p\delta_1 + (1 - p)\delta_0$ denotes the Bernoulli distribution with parameter p and δ_t is the Dirac measure for $\{t\}$.

Now consider perturbations G_{nu}, π_{nw} and Q_{nv} of G, π and Q, respectively, that are *Hellinger differentiable* in the following sense:

$$\int_{[0,\,1]^m} \left\{ n^{1/2} \Big\{ dG_{nu}^{1/2}(x) - dG^{1/2}(x) \Big\} - \frac{1}{2} u(x) dG^{1/2}(x) \right\}^2 dx \to 0,$$

$$\int_{[0,\,1]^m} \int_{\{0,\,1\}} \left\{ n^{1/2} \Big\{ dB_{\pi_{nw}}^{1/2}(z \,|\, x) - dB_{\pi}^{1/2}(z \,|\, x) \Big\} - \frac{1}{2} \{ z - \pi(x) \} w(x) dB_{\pi}^{1/2}(z \,|\, x) \right\}^2 G(dx) \to 0,$$

$$\int_{[0,1]^m}\int_{-\infty}^{\infty}\Big\{n^{1/2}\Big\{dQ_{nv}^{1/2}(y\,|\,x)-dQ^{1/2}(y\,|\,x)\Big\}-\frac{1}{2}v(x,y)dQ^{1/2}(y\,|\,x)\Big\}^2 G_1(dx)\to 0,$$

writing G_1 for the conditional distribution of X given that $\delta=1$. The perturbed distribution functions G_{nu}, $B_{\pi_{nw}}$ and Q_{nv} must satisfy the original model constraints, which requires their Hellinger derivatives to be restricted to suitable function spaces: u belongs to $\mathcal{L}_{2,0}(G)$, i.e. $u\in\mathcal{L}_2(G)$ and $\int_{[0,1]^m}u(x)\,G(dx)=0$; w belongs to

$$\mathcal{L}_2(G_\pi)=\Big\{w\in\mathcal{L}_2(G)\,:\,\int_{[0,1]^m}w^2(x)\pi(x)\{1-\pi(x)\}\,G(dx)<\infty\Big\},$$

writing $G_\pi(dx)=\pi(x)\{1-\pi(x)\}G(dx)$; and v belongs to

$$\mathcal{V}_0=\Big\{v\in\mathcal{L}_2(Q\otimes G_1)\,:\,\int_{-\infty}^{\infty}v(x,y)\,Q(dy\,|\,x)=0\Big\}.$$

Note that models for G_1, π and Q will imply further restrictions on the perturbations in order to satisfy those model assumptions. This means that u, w and v must be further restricted to subspaces of $\mathcal{L}_{2,0}(G)$, $\mathcal{L}_2(G_\pi)$ and $\mathcal{V}_0$, respectively. Here no model assumptions on G and π have been introduced, but model equation (1.1) does present a structural constraint on the conditional distribution Q of Y given X. This implies that we only have to identify the appropriate subspace $\mathcal{V}$ of $\mathcal{V}_0$ to account for this additional structure.

Since the covariates and the errors are assumed to be independent, we may write the density function dQ of Q as $dQ(x,y)=f(y-r(x))$. Using this notation, the constraint on $v\in\mathcal{V}_0$ now states that

$$\int_{-\infty}^{\infty}v(x,y)f(y-r(x))dy=0.$$

In order to derive the explicit form of the function space $\mathcal{V}$, we introduce further respective perturbations s and t for the unknown functions f and r, and we can write

$$dQ_{nv}(x,y)=dQ_{nst}(x,y)=f_{ns}\big(y-r_{nt}(x)\big),$$

where $f_{ns}(z)=f(z)\{1+n^{-1/2}s(z)\}$ and $r_{nt}(x)=r(x)+n^{-1/2}t(x)$ with $s\in\mathcal{S}$ and $t\in\mathcal{L}_2(G_1)$. Our assumptions on model (1.1) require the errors to have mean zero and the perturbed error density f_{ns} must integrate to one. Hence, $\mathcal{S}$ takes the form

$$\mathcal{S}=\Big\{s\in\mathcal{L}_2(F)\,:\,\int_{-\infty}^{\infty}s(z)f(z)\,dz=0,\int_{-\infty}^{\infty}zs(z)f(z)\,dz=0\Big\}.$$

We can simply restrict the perturbation t to belong to $\mathcal{L}_2(G_1)$, which follows from the fact that we do not assume a parametric form for r.

With the appropriate spaces $\mathcal{S}$ and $\mathcal{L}_2(G_1)$ identified, we can specify the appropriate form of $\mathcal{V}$. In the following arguments we will write "$\doteq$" to denote asymptotic equivalence, i.e. equality up to an additive term of order $o_p(n^{-1/2})$. As in Müller (2009), who considers a *parametric* (nonlinear) regression function, a brief sketch gives

$$\begin{aligned}
f_{ns}\big(y-r_{nt}(x)\big)&=f\big(y-r_{nt}(x)\big)\Big\{1+n^{-1/2}s\big(y-r_{nt}(x)\big)\Big\}\\
&=f\big(y-r(x)-n^{-1/2}t(x)\big)\Big\{1+n^{-1/2}s\big(y-r(x)-n^{-1/2}t(x)\big)\Big\}\\
&\doteq f(y-r(x))\Big\{1+n^{-1/2}\big\{s(y-r(x))+\ell(y-r(x))t(x)\big\}\Big\}.
\end{aligned}$$

Hence, we can write

$$dQ_{nst}(x,y)\doteq f(y-r(x))\Big\{1+n^{-1/2}\big\{s(y-r(x))+\ell(y-r(x))t(x)\big\}\Big\} \tag{2.1}$$

Equation (2.1) implies that $\mathcal{V}$ has the form

$$\mathcal{V}=\Big\{v(x,y)=s(y-r(x))+\ell(y-r(x))t(x)\,:\,s\in\mathcal{S}\text{ and }t\in\mathcal{L}_2(G_1)\Big\}.$$

We can see that $\dot{\Theta}$ is the set containing all possible Hellinger perturbations of the statistical model parameters:

$$\dot{\Theta} = \mathcal{L}_{2,0}(G) \times \mathcal{S} \times \mathcal{L}_2(G_1) \times \mathcal{L}_2(G_\pi).$$

The perturbed distribution $P_{n\gamma}$, with $\gamma = (u, s, t, w)$ in $\dot{\Theta}$, of the observation $(X, \delta Y, \delta)$ can be written

$$P_{n\gamma}(dx, dy, dz) \doteq G_{nu}(dx) B_{\pi_{nw(x)}}(dz \,|\, x)\Big\{ zQ_{nst}(dy \,|\, x) + (1-z)\delta_0(dy)\Big\}.$$

It then follows that $P_{n\gamma}$ is Hellinger differentiable with perturbation function

$$d_\gamma(x, y, z) = u(x) + z\big\{ s(y - r(x)) + \ell(y - r(x))t(x)\big\} + \{z - \pi(x)\}w(x), \tag{2.2}$$

and we have the stochastic expansion, writing dP for the density function of P,

$$\sum_{j=1}^n \log\left(\frac{dP_{n\gamma}(X_j, \delta_j Y_j, \delta_j)}{dP(X_j, \delta_j Y_j, \delta_j)}\right) = n^{-1/2}\sum_{j=1}^n d_\gamma(X_j, \delta_j Y_j, \delta_j) - \frac{1}{2}E\big[d_\gamma^2(X, \delta Y, \delta)\big] + o_p(1).$$

Since $n^{-1/2}\sum_{j=1}^n d_\gamma(X_j, \delta_j Y_j, \delta_j)$ is asymptotically normally distributed with mean zero and variance $E[d_\gamma^2(X, \delta Y, \delta)]$ by the central limit theorem, it follows for the expansion above to characterize local asymptotic normality in the present situation.

The efficient influence function of a differentiable functional is characterized by its canonical gradient, which takes the form $d_\gamma^*(X, \delta Y, \delta)$ for some $\gamma^* \in \dot{\Theta}$. This gradient is defined as the orthogonal projection of the gradient for the functional $E[h(\varepsilon)]$ (to be specified later) onto the tangent space given by the perturbed distributions $P_{n\gamma}$. It then follows from (2.2) for the tangent space $\mathscr{T}$ to be equal to the closure of the linear subspace formed by d_γ. Since d_γ is a sum of orthogonal elements we can write

$$\mathscr{T} = \Big\{u(X) \,:\, u \in \mathcal{L}_{2,0}(G)\Big\} \oplus \Big\{\delta v(X, Y) \,:\, v \in \mathcal{V}\Big\} \oplus \Big\{\{\delta - \pi(X)\}w(X) \,:\, w \in \mathcal{L}_2(G_\pi)\Big\}.$$

We are interested in the linear functional $E[h(\varepsilon)]$. In order to specify a gradient for $E[h(\varepsilon)]$, we need the directional derivative $\gamma_h \in \dot{\Theta}$ of $E[h(\varepsilon)]$, which is characterized by a limit as follows. As in Müller, Schick and Wefelmeyer (2004) we have, for every $s \in \mathcal{S}$,

$$\lim_{n\to\infty} n^{1/2}\left\{\int_{-\infty}^\infty h(z) f_{ns}(z)\, dz - \int_{-\infty}^\infty h(z) f(z)\, dz\right\} = E\big[h(\varepsilon)s(\varepsilon)\big] = E\big[h_0(\varepsilon)s(\varepsilon)\big],$$

where h_0 is the projection of h onto $\mathcal{S}$:

$$h_0(z) = h(z) - E\big[h(\varepsilon)\big] - \frac{z}{\sigma^2}E\big[\varepsilon h(\varepsilon)\big]. \tag{2.3}$$

Here σ^2 denotes the error variance. Hence, $E[h(\varepsilon)]$ is directionally differentiable, and (2.3) implies this directional derivative is $\gamma_h = (0, h_0, 0, 0)$. It then follows for $E[h(\varepsilon)]$ to have the gradient $h_0(\varepsilon)$, with h_0 given by (2.3).

By the convolution theorem (see, for example, Section 2 of Schick, 1993), the unique canonical gradient $g^*(X, \delta Y, \delta)$ is obtained by orthogonally projecting the gradient $h_0(\varepsilon)$ of $E[h(\varepsilon)]$ onto the tangent space $\mathscr{T}$. Hence, $g^*(X, \delta Y, \delta)$ must be of the form

$$g^*(X, \delta Y, \delta) = u^*(X) + \delta\big\{s^*(\varepsilon) + \ell(\varepsilon)t^*(X)\big\} + \big\{\delta - \pi(X)\big\}w^*(X), \tag{2.4}$$

which satisfies

$$E\big[h_0(\varepsilon)s(\varepsilon)\big] = E\big[g^*(X, \delta Y, \delta)d_\gamma(X, \delta Y, \delta)\big] \tag{2.5}$$

for every $\gamma \in \dot{\Theta}$. A straightforward calculation shows the right-hand side of (2.5) is equal to

$$\begin{aligned} &E\big[g^*(X, \delta Y, \delta)d_\gamma(X, \delta Y, \delta)\big] \\ &= E\big[u^*(X)u(X)\big] + E\big[\delta\big\{s^*(\varepsilon) + \ell(\varepsilon)t(X)\big\}\big\{s(\varepsilon) + \ell(\varepsilon)t^*(X)\big\}\big] \\ &\quad + E\big[\{\delta - \pi(X)\}^2 w^*(X)w(X)\big] \end{aligned}$$

$$= E\big[u^*(X)u(X)\big] + E\delta\Big\{E\big[s^*(\varepsilon)s(\varepsilon)\big] + E\big[\ell_0(\varepsilon)s^*(\varepsilon)\big]E_1\big[t(X)\big] + E\big[\ell_0(\varepsilon)s(\varepsilon)\big]E_1\big[t^*(X)\big] + JE_1\big[t^*(X)t(X)\big]\Big\} + E\big[\pi(X)\{1-\pi(X)\}w^*(X)w(X)\big],$$

where J is the Fisher information given in Assumption 3 and $\ell_0(\varepsilon)$ is the projection of $\ell(\varepsilon)$ onto $\mathcal{V}$, i.e. $\ell_0(\varepsilon) = \ell(\varepsilon) - (\varepsilon/\sigma^2)$. The notation E_1 indicates the expectation is with respect to the conditional distribution G_1. For convenience, we introduce the quantity J_0 which is calculated analogously to J:

$$J_0 = E\big[\ell_0^2(\varepsilon)\big] = E\left[\left\{\ell(\varepsilon) - \frac{\varepsilon}{\sigma^2}\right\}^2\right] = J - \frac{1}{\sigma^2}.$$

With appropriate choices of γ in $\dot{\Theta}$, it easily follows from (2.5) for $u^* = w^* = 0$. Now choosing the zero function for the function u, (2.5) becomes

$$E\big[h_0(\varepsilon)s(\varepsilon)\big] = E\delta\Big\{E\big[s(\varepsilon)s^*(\varepsilon)\big] + E\big[\ell_0(\varepsilon)s(\varepsilon)\big]E_1\big[t^*(X)\big]\Big\},$$

which must hold for all $s \in \mathcal{S}$. This implies

$$s^*(z) = \frac{h_0(z)}{E\delta} - \ell_0(z)E_1\big[t^*(X)\big], \qquad z \in \mathbb{R}. \tag{2.6}$$

Following Müller, Schick and Wefelmeyer (2004), we can consider $\mathcal{L}_2(G_1)$ written as an orthogonal sum of functions with mean zero and of constants, i.e. we write $\mathcal{L}_2(G_1) = \mathcal{L}_{2,0}(G_1) \oplus [1]$. This means we can decompose $t(X)$ into $t(X) = \{t(X) - E_1[t(X)]\} + E_1[t(X)]$. Finally, choosing the zero function for s and inserting s^* from (2.6), (2.5) becomes

$$\begin{aligned} 0 &= \frac{1}{E\delta}E\big[\ell_0(\varepsilon)h_0(\varepsilon)\big]E_1\big[t(X)\big] + (J - J_0)E_1\big[t^*(X)\big]E_1\big[t(X)\big] \\ &\quad + JE_1\left[\Big\{t^*(X) - E_1\big[t^*(X)\big]\Big\}\Big\{t(X) - E_1\big[t(X)\big]\Big\}\right] \\ &= \left\{\frac{1}{E\delta}E\big[\ell_0(\varepsilon)h_0(\varepsilon)\big] + \frac{1}{\sigma^2}E_1\big[t^*(X)\big]\right\}E_1\big[t(X)\big] \\ &\quad + JE_1\left[\Big\{t^*(X) - E_1\big[t^*(X)\big]\Big\}\Big\{t(X) - E_1\big[t(X)\big]\Big\}\right], \end{aligned}$$

which must hold for every $t \in \mathcal{L}_2(G_1)$. This implies t^* is equal to its mean $E_1[t^*(X)]$ and

$$E_1\big[t^*(X)\big] = -\frac{\sigma^2}{E\delta}E\big[h_0(\varepsilon)\ell_0(\varepsilon)\big]. \tag{2.7}$$

Therefore, combining (2.7) with the fact that t^* must be equal to its mean yields $t^*(x) = -\sigma^2(E\delta)^{-1}E[h_0(\varepsilon)\ell_0(\varepsilon)]$, $x \in [0,1]^m$. Combining the fact that $s^* = w^* = 0$ with (2.6) and (2.7), we obtain the following result:

LEMMA 1. *The canonical gradient of $E[h(\varepsilon)]$ is $g^*(X, \delta Y, \delta)$ given in* (2.4) *that characterized by $(0, s^*, t^*, 0)$, with*

$$s^*(z) = \frac{1}{E\delta}\Big\{h_0(z) + \sigma^2 E\big[h_0(\varepsilon)\ell_0(\varepsilon)\big]\ell_0(z)\Big\}, \qquad z \in \mathbb{R},$$

$$t^*(x) = -\frac{\sigma^2}{E\delta}E\big[h_0(\varepsilon)\ell_0(\varepsilon)\big], \qquad x \in [0,1]^m,$$

where $\sigma^2 = E[\varepsilon^2]$ is the error variance, h_0 is given in (2.3) *and $\ell_0(\varepsilon) = \ell(\varepsilon) - (\varepsilon/\sigma^2)$.*

An estimator $\hat{\mu}$ of $E[h(\varepsilon)]$ is called *efficient*, in the sense of Hájek and Le Cam, when $\hat{\mu}$ is asymptotically linear with influence function equal to the canonical gradient $g^*(X, \delta Y, \delta)$ that

characterizes $E[h(\varepsilon)]$, i.e. if the expansion holds:

$$n^{1/2}\Big\{\hat{\mu} - E\big[h(\varepsilon)\big]\Big\} = n^{-1/2}\sum_{j=1}^{n} g^*(X_j, \delta_j Y_j, \delta_j) + o_p(1).$$

A straightforward calculation combining the result of Lemma 1 with the display above and formula (2.4) yields:

COROLLARY 1. *Consider the nonparametric regression model with responses missing at random. An efficient estimator $\hat{\mu}$ of $E[h(\varepsilon)]$ must satisfy the expansion*

$$n^{1/2}\Big\{\hat{\mu} - E\big[h(\varepsilon)\big]\Big\} = n^{-1/2}\sum_{j=1}^{n} \frac{\delta_j}{E\delta}\Big\{h(\varepsilon_j) - E\big[h(\varepsilon)\big] - E\big[\ell(\varepsilon)h(\varepsilon)\big]\varepsilon_j\Big\} + o_p(1).$$

REMARK 2. Müller, Schick and Wefelmeyer (2004) construct residual-based estimators $n^{-1}\sum_{j=1}^{n} h(\hat{\varepsilon}_j)$ for estimating $E[h(\varepsilon)]$ in the full model. In their Section 2, they give conditions for the i.i.d. representation:

$$n^{-1/2}\sum_{j=1}^{n} h(\hat{\varepsilon}_j) = n^{-1/2}\sum_{j=1}^{n}\Big\{h(\varepsilon_j) - E\big[h'(\varepsilon)\big]\varepsilon_j\Big\} + o_p(1),$$

which characterizes an efficient estimator. (For simplicity, we assume in this remark that h is differentiable.) Note that $E[h'(\varepsilon)] = E[\ell(\varepsilon)h(\varepsilon)]$. Hence, using the transfer principle, we see that the complete case versions of their estimators have the expansion from the previous corollary, and, therefore, these estimates are also efficient in the MAR model.

The function $h(\varepsilon) = \mathbf{1}[\varepsilon \le t]$ is of particular interest because many statistical methods are residual-based and require estimation of the error distribution function. Using Corollary 1 with this particular $h(\varepsilon)$, we obtain an expansion for the residual-based empirical distribution function:

COROLLARY 2. *Consider the nonparametric regression model with responses missing at random. An estimator $\hat{F}$ of the error distribution function F is efficient, if it satisfies the expansion*

$$n^{1/2}\Big\{\hat{F}(t) - F(t)\Big\} = n^{-1/2}\sum_{j=1}^{n} \frac{\delta_j}{E\delta}\Big\{\mathbf{1}\big[\varepsilon_j \le t\big] - F(t) + f(t)\varepsilon_j\Big\} + o_p(1).$$

Note, this is the expansion of the complete case estimator $\hat{\mathbb{F}}_c$ from the previous section, which provide the proof of the second assertion in Theorem 2.

3. Simulation results

To conclude the article, we present a brief simulation study of the previous results in two important examples. The first example compares the efficiency property of $\hat{\mathbb{F}}_c$, which is constructed using only the completely observed data, to another estimator $\tilde{\mathbb{F}}$, which is constructed using an imputation methodology. In the second example, we consider applying a goodness-of-fit test for normal errors to the residuals from the nonparametric regression constructed by only the complete cases and by the same imputation methodology that was implemented in the first example. In both examples, we assume a nonparametric regression model (1.1), but choosing

$$r(x) = x^3 - x^2 + x + \cos\big((3\pi/2)x\big),$$

which we expect preserves the nonparametric nature of the studies.

The covariates are generated from a uniform distribution and the errors from a normal distribution: $X_j \sim U(-1,1)$ and $\varepsilon_j \sim N(0,1)$ for $j = 1, \ldots, n$; see Figure 1 for a scatterplot of a typical simulated dataset. Finally, the indicators δ_j have a Bernoulli distribution with

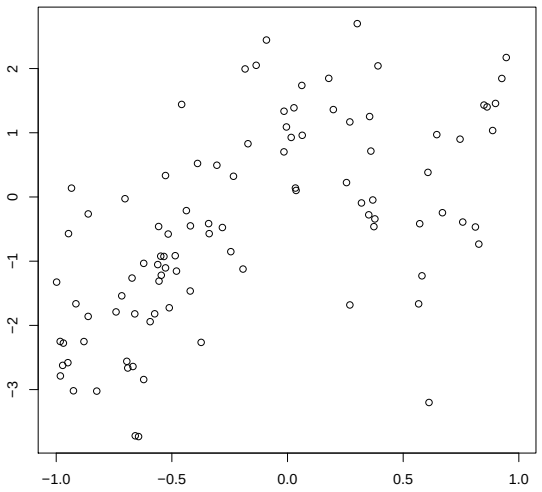

FIGURE 1. *A scatter plot of* 100 *data points from a typical simulated dataset.*

proportion function parameter $\pi(x) = P(\delta = 1|X = x)$, which is chosen to be the logistic distribution function with a mean of zero and scale parameter of one:

$$\pi(x) = \frac{1}{1 + e^{-x}}.$$

Consequently, the average amount of missing data is 50% and ranges between 27% and 73%. Finally, we work with the local linear smoother (it is easy to see that r is Lipschitz and, therefore, it belongs to the Hölder space $H(1, 1)$), and the bandwidth sequence $\{c_n\}_{n\geq 1}$ is taken as $c_n = 1.25\{n \log(n)\}^{-1/4}$. The assumptions of Theorem 2 are then satisfied.

3.1. Example 1: Simulation of asymptotic mean squared error. We consider two estimators of the error distribution function. The first estimator is the proposed complete case estimator $\hat{\mathbb{F}}_c$ and the second is a 'tuned' version of $\hat{\mathbb{F}}_c$ that utilizes an imputation technique. Similar to González-Manteiga and Pérez-González (2006), we take the initial local polynomial complete case estimator $\hat{r}_c$ (see equation (1.4)) to produce the completed sample $(X_j, \hat{Y}_j)$, $j = 1, \ldots, n$. We chose to fully impute the responses, i.e. $\hat{Y}_j = \hat{r}_c(X_j)$. This is a variation of the approach of González-Manteiga and Pérez-González (2006), who work with partially imputed responses $\hat{Y}_j = \delta_j Y_j + (1 - \delta_j)\hat{r}_c(X_i)$. A new local polynomial estimator $\hat{r}^*(\cdot)$ is then constructed from the completed sample. When Y is observed, we can compute adjusted residuals $\hat{\varepsilon}^*_j = Y_j - \hat{r}^*(X_j)$ based on the updated estimator $\hat{r}^*$. This means we still work with complete cases when estimating F, but now the estimated regression has changed.

Using these residuals we obtain the new tuned estimator

$$\tilde{\mathbb{F}}(t) = \frac{1}{N} \sum_{j=1}^{n} \delta_j \mathbf{1}\big[\hat{\varepsilon}^*_j \leq t\big].$$

The results in the previous sections show the complete case estimator $\hat{\mathbb{F}}_c$ is an (asymptotically) efficient estimator of F. Our discussion in Remark 1 also suggests the tuned estimator $\tilde{\mathbb{F}}$ is also efficient; i.e. both estimators are asymptotically equivalent. We expect that $\tilde{\mathbb{F}}$ can be expanded in the same way as $\hat{\mathbb{F}}_c$:

$$\sup_{t\in\mathbb{R}} \left| \frac{1}{n} \sum_{j=1}^{n} \frac{\delta_j}{E\delta} \Big\{ \mathbf{1}\big[\hat{\varepsilon}^*_j \leq t\big] - \mathbf{1}\big[\varepsilon_j \leq t\big] \Big\} - f(t) \int_{[0,1]^m} \hat{a}^*(x)\, G_1(dx) \right| = o_p(n^{-1/2}),$$

n \ t	-1.5		-1		0		1		1.5	
50	0.1141	0.0987	0.2705	0.2087	0.1702	0.1884	0.2865	0.2220	0.1179	0.1009
250	0.1018	0.0930	0.1800	0.1634	0.2021	0.2071	0.2022	0.1972	0.1201	0.1165
1000	0.0991	0.0945	0.1668	0.1625	0.1865	0.1997	0.1706	0.1780	0.1000	0.1008
10000	0.0925	0.0920	0.1567	0.1537	0.2068	0.2274	0.1690	0.1752	0.0953	0.0975
true	0.0911	–	0.1498	–	0.1816	–	0.1498	–	0.0911	–

TABLE 1. *Simulated and true asymptotic MSE of* $n^{1/2}\{\hat{\mathbb{F}}_c - F\}$ *and* $n^{1/2}\{\tilde{\mathbb{F}} - F\}$ *at the points* -1.5, -1, 0, 1 *and* 1.5.

where $\hat{a}^*(x)$ is now an approximation of the difference $\hat{r}^*(x) - r(x)$ (cf. equation (1.5) in Remark 1). The integral in the display above can be written as

$$\int_{[0,1]^m} \hat{a}^*(x)\, G_1(dx) = \int_{[0,1]^m} \hat{a}_c(x)\, G_1(dx) + \int_{[0,1]^m} \big\{\hat{a}^*(x) - \hat{a}_c(x)\big\}\, G_1(dx).$$

Since $\hat{a}^*(x) - \hat{a}_c(x)$ approximates the difference $\hat{r}^*(x) - \hat{r}_c(x)$ of two consistent estimators of $r(x)$, we expect the last term in the display above to be asymptotically negligible. Repeating the arguments from Remark 1 would then give the desired expansion:

$$\sup_{t\in\mathbb{R}} \left| \frac{1}{N} \sum_{j=1}^{n} \delta_j \Big\{ \mathbf{1}\big[\hat{\varepsilon}_j^* \le t\big] - \mathbf{1}\big[\varepsilon_j \le t\big] - \varepsilon_j f(t) \Big\} \right| = o_p(n^{-1/2}),$$

and, hence, both $\hat{\mathbb{F}}_c$ and $\tilde{\mathbb{F}}$ should have the same asymptotic expansion, i.e. both estimators are asymptotically equivalent.

In order to further check the conjecture that both estimators are asymptotically equivalent, we conducted a simulation study using 1000 trials. We considered four sample sizes and five different points at which the error distribution function was evaluated. The findings are summarized in Table 1. Note, we also implemented another estimator, which uses partial imputation to complete the sample as suggested by González-Manteiga and Pérez-González (2006), but our approach performed slightly better and so we only report the results for our tuned estimator $\tilde{\mathbb{F}}$. For the second smoothing step we chose the same bandwidth as in the first step, $c_n = 1.25\{n\log(n)\}^{-1/4}$.

These results show the simulated MSE (multiplied by n) of the efficient estimator $\hat{\mathbb{F}}_c$ is very close to the true asymptotic MSE (which equals the asymptotic variance and can be calculated using the results of Theorem 2). We can also see the asymptotic MSE estimates of $\tilde{\mathbb{F}}$ are similar to those of $\hat{\mathbb{F}}_c$ at large sample sizes. This provides further evidence for the two approaches to be asymptotically equivalent as conjectured. The simulated MSE values of $\tilde{\mathbb{F}}$, however, more closely match the true asymptotic MSE values at low sample sizes, which we expect is due to the imputation technique. However, at the point 0 both estimators perform very similarly for all sample sizes. A possible explanation of this behavior is the point 0 is also the median of this error distribution, and we believe the imputation technique is least helpful in this case from our discussion above conjecturing on the expansion of $\tilde{\mathbb{F}}$.

3.2. Example 2: Simulating a goodness-of-fit test for normal errors. We now consider a test proposed by Müller, Schick and Wefelmeyer (2012) for the full model with multivariate covariates. This test was also examined by Koul, Müller and Schick (2012) in the MAR model with a one-dimensional covariate, but without simulations. Both articles study versions of a martingale transform test developed by Khmaladze and Koul (2009). Under the null hypothesis, these test statistics have limiting distributions given by $\sup_{0<t\le 1} |B(t)|$, where $B(t)$ is the standard Brownian motion. These test statistics are asymptotically distribution free because the limiting distribution does not depend on any unknown parameters, which would

	$N(0,2)$		χ^2_1-1		t_4		Laplace$(0,2)$	
n	T_c	T_ι	T_c	T_ι	T_c	T_ι	T_c	T_ι
50	0.022	0.025	0.489	0.535	0.099	0.108	0.095	0.119
200	0.030	0.028	1.000	1.000	0.457	0.463	0.459	0.483

TABLE 2. Test for normally distributed errors. Simulated level is given by $N(0,2)$ figures.

have to be estimated. This is very useful because the corresponding complete case statistics have the same limiting distributions in this case, which is a consequence of the transfer principle. Hence, the decision rule remains unchanged in the MAR model. For example, setting the level of the test to 0.05, we reject H_0 when the test statistic exceeds 2.2414, which is the upper 5% quantile of the distribution of $\sup_{0<t\leq 1}|B(t)|$.

Writing ϕ for the density function of the standard normal distribution and σ^2 for the error variance, the null hypothesis of normal errors is

$$H_0 : \exists\, \sigma > 0 \,:\, f(t) = \frac{1}{\sigma}\phi\left(\frac{t}{\sigma}\right), \quad t \in \mathbb{R}.$$

In order to introduce the test statistic T_c, define $h(x) = (1, -\phi'(x)/\phi(x), -(x\phi(x))'/\phi(x))^T$ and

$$H(t) = \int_{-\infty}^{t} h^T(t)\Gamma^{-1}(t)\phi(t)\,dt,$$

where $\Gamma(t) = \int_t^\infty h(u)h^T(u)\phi(u)\,du$ (see Müller, Schick and Wefelmeyer, 2012, and Koul, Müller and Schick, 2012, for an explicit definition of Γ and for more details). Following Koul, Müller and Schick (2012), we have the test statistic

$$T_c = \sup_{t\in\mathbb{R}} \left| N^{-1/2} \sum_{j=1}^{n} \delta_j \Big\{ \mathbf{1}\big[\hat{Z}_{j,c} \leq t\big] - H\big(t \wedge \hat{Z}_{j,c}\big) h\big(\hat{Z}_{j,c}\big) \Big\} \right|.$$

Note, this statistic is based on our proposed estimator $\hat{\mathbb{F}}_c$, but with *scaled* residuals $\hat{Z}_{j,c} = \hat{\varepsilon}_{j,c}/\hat{\sigma}_c$, where $\hat{\sigma}_c$ is the complete case version of the residual-based empirical estimator, i.e. $\hat{\sigma}_c = \sqrt{\hat{\sigma}_c^2}$ with

$$\hat{\sigma}_c^2 = \frac{1}{N}\sum_{j=1}^{n} \delta_j \hat{\varepsilon}_{j,c}^2 = \frac{1}{N}\sum_{j=1}^{n} \delta_j \big\{Y_j - \hat{r}_c(X_j)\big\}^2.$$

Recall that, under the MAR assumption, ε and δ are independent. Hence, $\hat{\sigma}_c^2$ is a consistent estimator of $\mathrm{Var}(\varepsilon|\delta = 1) = \mathrm{Var}(\varepsilon) = \sigma^2$.

We are interested in studying the performance of T_c in the MAR model, and we wish to compare it with the corresponding statistic T_ι that is based on the tuned estimator $\tilde{\mathbb{F}}$. Here T_ι has exactly the same form as T_c but with all $\hat{\varepsilon}_{j,c}$ replaced by the adjusted residuals $\hat{\varepsilon}_j^* = Y_j - \hat{r}^*(X_j)$. For the simulations, we consider the same scenario as in the previous example, but now also admit some other models for the error distribution.

First we look at the $N(0,2)$ distribution to allow verification of the (5%) level of the test. To check the power of the test, we generated errors from a mean shifted $\chi^2(1)$ distribution, a $t(4)$ distribution and a Laplace distribution with mean 0 and variance 2. The simulation study is based on 1000 runs and samples of size 50 and 200.

Table 2 shows, when the errors are normally distributed (and the null hypothesis is true), the test using T_c rejects the null hypothesis 2.2% of the time for samples of size 50, and 3% of the time for samples of size 200. This indicates the test using T_c is slightly conservative. We find similar conservative behavior in the test using T_ι, where the hypothesis of normality is rejected 2.5% and 2.8% of the time for sample sizes 50 and 200, respectively. When the null hypothesis is not true, the power figures are fairly close for both tests. The test using T_ι seems to be more powerful for low sample sizes, which is expected from the results of the first example. The differences are less pronounced for the larger sample size of 200, suggesting

that the two approaches are asymptotically equivalent – which is also what we would expect given the discussion and the simulation results in the previous example. Summing up, both test procedures have similar performance. The test based on T_c appears to be the better choice for moderately large (or large) samples because it is easier to implement.

Acknowledgements

Ursula U. Müller was supported by NSF Grant DMS 0907014. The authors thank the referees for a number of suggestions that improved the manuscript. We also thank Susan Davis for her helpful comments on an earlier draft.

References

[1] Efromovich, S. (2011a). Nonparametric regression with responses missing at random. *J. Statist. Plann. Inference* **141**, 3744-3752.
[2] Efromovich, S. (2011b). Nonparametric regression with predictors missing at random. *J. Amer. Statist. Assoc.* **106**, 306-319.
[3] Khmaladze, E.V. and Koul, H.L. (2004). Martingale transforms goodness-of-fit tests in regression models. *Ann. Statist.* **32**, 995-1034.
[4] Khmaladze, E.V. and Koul, H.L. (2009). Goodness-of-fit problem for errors in non-parametric regression: distribution free approach. *Ann. Statist.* **37**, 3165-3485.
[5] Koul, H.L., Müller, U.U. and Schick, A. (2012). The transfer principle: a tool for complete case analysis. *Ann. Statist.* **40**, 3031-3049.
[6] González-Manteiga, W. and Pérez-González, A. (2006). Goodness-of-fit tests for linear regression models with missing response data. *Canad. J. Statist.* **34**, 149-170.
[7] Liang, H., Wang, S. and Carroll, R. (2007). Partially linear models with missing response variables and error-prone covariates. *Biometrika* **94**, 185-198.
[8] Little, R.J.A. and Rubin, D.B. (2002). *Statistical analysis with missing data*. Second edition, Wiley-Interscience.
[9] Molenberghs, G. and Kenward, M. (2007). *Missing Data in Clinical Studies*. Wiley.
[10] Müller, U.U. (2009). Estimating linear functionals in nonlinear regression with responses missing at random. *Ann. Statist.* **37**, 2245-2277.
[11] Müller, U.U., Schick, A. and Wefelmeyer, W. (2004). Estimating linear functionals of the error distribution in nonparametric regression. *J. Statist. Plann. Inference* **119**, 75-93.
[12] Müller, U.U., Schick, A. and Wefelmeyer, W. (2006). Imputing responses that are not missing. In *Probability, Statistics and Modelling in Public Health*, eds. Nikulin, M., Commenges, D. and Huber, C., Springer, 350-363.
[13] Müller, U.U., Schick A. and Wefelmeyer W. (2007). Estimating the error distribution in semiparametric regression. *Statist. Decisions* **25**, 1-18.
[14] Müller, U.U., Schick, A. and Wefelmeyer, W. (2009). Estimating the error distribution function in nonparametric regression with multivariate covariates. *Statist. Probab. Lett.* **79**, 957-964.
[15] Müller, U.U., Schick, A. and Wefelmeyer, W. (2012). Estimating the error distribution function in semiparametric additive regression models. *J. Statist. Plann. Inference* **142**, 552-566.
[16] Neumeyer, N. and Van Keilegom, I. (2010). Estimating the error distribution in nonparametric multiple regression with applications to model testing. *J. Multivariate Anal.* **101**, 1067-1078.
[17] Schick, A. (1993). On efficient estimation in regression models. *Ann. Statist.* **21**, 1486-1521.
[18] Stute, W. (1997). Nonparametric model checks for regression. *Ann. Statist.* **25**, 613-641.
[19] Tsiatis, A. (2006). *Semiparametric Theory and Missing Data*. Springer.